arXiv:1612.06096v1 [cs.CV] 19 Dec 2016

X-ray In-Depth Decomposition: Can Deep Learning Reveal The Latent Structures?

Shadi Albarqouni[1], Javad Fotouhi[2], and Nassir Navab[1,2]

[1] Computer Aided Medical Procedures, Technische Universität München, Munich, Germany
[2] Computer Aided Medical Procedures, Johns Hopkins University, Baltimore, MD, Unites States.
shadi.albarqouni@tum.de

Abstract. X-ray radiography is the most readily available imaging modality and has a broad range of applications that spans from diagnosis to intra-operative guidance in cardiac, orthopedics, and trauma procedures. Proper interpretation of the hidden and obscured anatomy in X-ray images remains a challenge and often requires high radiation dose and imaging from several perspectives.
In this work, we aim at decomposing the conventional X-ray image into d X-ray components of independent, non-overlapped, clipped sub-volumes using deep learning approach. Despite the challenging aspects of modelling such a highly ill-posed problem, exciting and encouraging results are obtained paving the path for further contributions in this direction.

1 Introduction

Two-dimensional X-ray is the primary imaging modality for clinical and interventional radiology. After several decades of its deployment, still proper determination between the healthy and abnormal anatomy is challenging and heavily relies on observer's perceptual skills and judgment [8,9,3]. Moreover, the accumulation of X-ray attenuation displayed on the X-ray image leads to anatomy obscuration. In other words, a single pixel in an X-ray image may partially contribute to soft tissue as well as hard tissue. In chest X-ray imaging, interpretation difficulties arise when cardiac silhouette takes place and makes the determination between anatomy challenging. In these images, observing and distinguishing the anterior and posterior ribs, sternum, spine, and the vasculature requires skilled clinicians.

In the surgical setting, a C-arm X-ray imaging device is used to acquire intra-operative fluoroscopic images. To better understand the underlying anatomy and navigate in 3D, the C-arm is frequently repositioned and images from multiple perspectives are acquired. This procedure is mentally and physically challenging for the surgical team and results in their frustration. Moreover, additional effort is required to visualize the vasculature when the contrast among the soft tissue is low. To this end, phase contrast X-ray techniques are introduced that yield high contrast for soft tissue imaging. These methods are based on measuring the phase

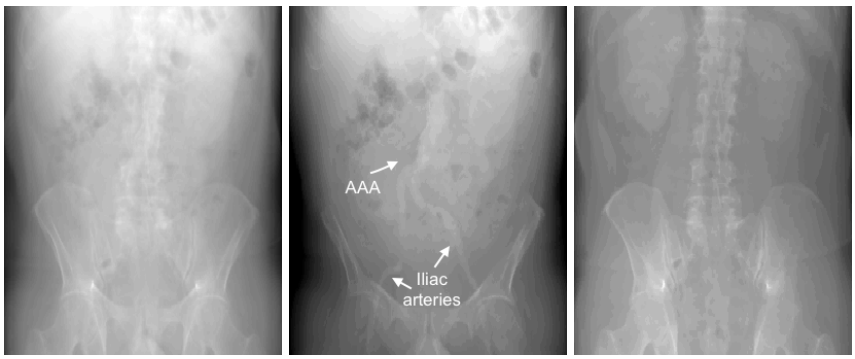

(a) (b) (c)

Fig. 1. Simulated X-ray images for (a) abdominal: $\boldsymbol{\mu_T}$, (b) AAA sub-volume: $\boldsymbol{\mu_A}$, and (c) spine sub-volume: $\boldsymbol{\mu_S}$. Region of interests (white arrows) are quite visible.

shift and X-ray diffraction after hitting the tissue [2]. An alternative solution is to use dual energy X-ray absorption, where X-ray images are acquired at two different energy levels, and later subtracted to only visualize the anatomy of a particular absorption level [5]. The primary application of this technique is the estimation of body composition, particularly the bone mineral density. Existing X-ray angiography solutions are mainly based on contrast-enhanced imaging, where a pre-contrast image is subtracted from a contrast-enhanced image to remove the bony anatomy [7]. Neither of the aforementioned methods provide a distinction between the anatomy at different depth and yet distinguishing the layers of the anterior and posterior anatomy depends on the surgeon's experience and judgment.

To improve the perceptual problem in 2D X-ray radiographs, prior knowledge from pre-operative patient CT data is used to restore depth data [15]. However, accurate registration between 3D preoperative CT scan and 2D X-ray image is required. Estimating the relative depth of the X-ray image relative to the CT data allows to modulate X-ray intensities (cutting-off at a certain depth) and introducing additional depth cues into the X-ray image. In [12], an interactive virtual mirror is integrated into the 2D/3D fusion of X-ray and CT data to facilitate the localization of aneurysm. The virtual mirror generates desired perspectives of the 3D data and displays 2D and 3D data on a single display. To further improve the depth perception of X-ray acquisitions, colored depth maps are computed using ray casting on 3D data, and are later integrated into the interventional X-ray while preserving the original gray-scale level of the X-ray [13,14]. These methods rely on accurate 2D/3D registration of data which by itself is a complex problem when the anatomy in the pre-interventional X-ray is dissimilar to the pre-operative patient data due to deformations. Last but not least, access to patient CT data is not available for the majority of cases.

Recent advances in machine learning has lead to adoption of learning-based approaches to tackle challenging problems in medical image processing. Convolutional neural networks are proposed in [6] to efficiently register 2D and 3D medical data. Similar models are used to estimate the depth of the anatomy from a single-view X-ray radiograph [1].

In this work, we employ a novel learning-based approach that enables the decomposition of radiographs into distinct layers, and mitigates the aforementioned limitations. To the best of our knowledge, this is the very first attempt to perform in-depth decomposition of a single X-ray image. We believe this work will pave the path for further improvements that can contribute to better interpretation of complex anatomy in chest X-ray images to distinguish the margins of the diaphragm, rib cage, and spine, as well as for angiography and orthopedics interventions.

2 Methodology

In this work, we aim at decomposing the conventional X-ray image into d X-ray components of independent, non-overlapped, clipped sub-volumes (c.f. Fig 1). Our methodology, depicted in Fig. 3, utilizes pre-operative CT scan, which is commonly acquired two weeks prior to the intervention, where the clipping planes are defined by the clinicians. To generate the training data, simulated X-ray images, so-called Digitally Reconstructed Radiographs (DRRs), are generated along trajectories which cover the region of interest (ROI), *i.e.* Abdominal Aortic Aneurysm (AAA). First, the CT volume is sliced in the anterior-posterior plane to several layers. DRRs are generated from different perspectives of the sub-volumes as well as the original CT data. These generated DRRs form the training set of our model. In the testing phase, for any given X-ray image, decomposition of d X-ray components is obtained.

2.1 Beer-Lambert Law

The intensity I of a given X-ray image is formed using the Beer-Lambert law by accumulating the Hounsfield Units (HU) of the CT volume along the ray,

$$I = I_0 * e^{-\int \mu(x)\, dx} \tag{1}$$

where $\mu(\cdot)$: $\mathbb{R}^3 \rightarrow \mathbb{R}^1$ is the X-ray absorption function, and I_0 is the initial radiation. Under the assumption of mono-energetic function, the intensity of X-ray for a clipped volume at x_c, is formulated as:

$$I = I_0 * e^{-(\int^{x_c} \mu(x)\, dx + \int_{x_c} \mu(x)\, dx)} \tag{2}$$

This equation is simplified by assuming the maximal radiation, *i.e* $I_0 = 1$, and log-normalizing both sides:

$$-\log \frac{I}{I_0} = \int^{x_c} \mu(x)\, dx + \int_{x_c} \mu(x)\, dx \tag{3}$$

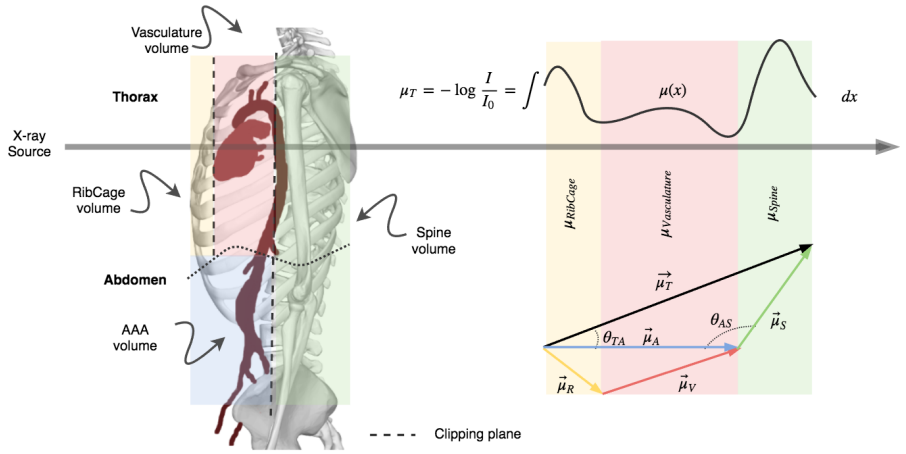

Fig. 2. Illustrative diagram of clipped CT volume (left) and X-ray decomposition in vector space (right).

By substituting Eq. 1 into Eq. 3, we obtain

$$\int^{x_c} \mu(x)\,dx = \int \mu(x)\,dx - \int_{x_c} \mu(x)\,dx, \tag{4}$$

which can be generalized for an arbitrary clipped volume between x_a and x_b:

$$\int_{x_a}^{x_b} \mu(x)\,dx = \int \mu(x)\,dx - \left(\int^{x_a} \mu(x)\,dx + \int_{x_b} \mu(x)\,dx\right) \tag{5}$$

In other words, the X-ray intensity of any clipped volume can be computed using a linear composition of the original X-ray image and the clipped X-ray images. For discrete signals, Eq. 5 is written as:

$$\mu_{ab} = \sum_{x_a}^{x_b} \mu_x = \sum_{x_0}^{x_a} \mu_x - \left(\sum_{x_0}^{x_a} \mu_x + \sum_{x_b}^{x_N} \mu_x\right) \tag{6}$$

and shorten without loss of generality to

$$\boldsymbol{\mu_{ab}} = \boldsymbol{\mu_T} - (\boldsymbol{\mu_a} + \boldsymbol{\mu_b}) \tag{7}$$

where $\boldsymbol{\mu_{ab}}$ stands for the X-ray image of clipped volume between plane x_a and x_b as depicted in Fig. 2.

2.2 Problem Formulation

Our objective is to predict d independent X-ray components from the original X-ray image $\boldsymbol{\mu_T}$, where the sum of stacked d ground truth components $\boldsymbol{\mu_d} \in$

$\mathbb{R}^{h\times w\times d}$ is equal to the original image:

$$\boldsymbol{\mu_T} = \sum_d \boldsymbol{\mu_d}, \tag{8}$$

As an example, the original X-ray image $\boldsymbol{\mu_T}$ can be decomposed into two components ($d = 2$) as:

$$\boldsymbol{\mu_T} = \boldsymbol{\mu_A} + \boldsymbol{\mu_S}, \tag{9}$$

where $\boldsymbol{\mu_T}, \boldsymbol{\mu_A}$, and $\boldsymbol{\mu_S} \in \mathbb{R}^{h\times w}$ are the corresponding X-ray images of chest (thorax), AAA and spine sub-volumes respectively. Our objective is that for any given X-ray image $\boldsymbol{\mu_T}$, the proposed model $f(\cdot)$ is able to predict decomposed layers as:

$$\hat{\boldsymbol{\mu_d}} = \{\hat{\boldsymbol{\mu_A}}, \hat{\boldsymbol{\mu_S}}\} = f(\boldsymbol{\mu_T}; w), \quad s.t. \quad \sum_d \hat{\boldsymbol{\mu_d}} = \hat{\boldsymbol{\mu_A}} + \hat{\boldsymbol{\mu_S}} = \boldsymbol{\mu_T}, \tag{10}$$

where $\hat{\boldsymbol{\mu_d}}$ is the predicted output and w is the model parameters. This objective function can be formulated using Lagrnagian multiplier as:

$$\arg\min_{\hat{\boldsymbol{\mu_d}}} \underbrace{\mathcal{L}_d\left(\hat{\boldsymbol{\mu_d}}, \boldsymbol{\mu_d}\right)}_{decomposition} + \lambda_r \underbrace{\mathcal{L}_r(\hat{\boldsymbol{\mu_d}}, \boldsymbol{\mu_T})}_{reconstruction} \tag{11}$$

where λ_r is the regularization paramter, and $\mathcal{L}_d(\cdot,\cdot)$ is the eleastic-net loss [16] between the predicted decompositions and their corresponding ground-truths:

$$\mathcal{L}_d\left(\hat{\boldsymbol{\mu_d}}, \boldsymbol{\mu_d}\right) = \lambda_d \underbrace{\|\hat{\boldsymbol{\mu_d}} - \boldsymbol{\mu_d}\|_2}_{smoothness} + (1-\lambda_d) \underbrace{\|\hat{\boldsymbol{\mu_d}} - \boldsymbol{\mu_d}\|_1}_{sparsity}, \tag{12}$$

where λ_d controls the tradeoff between ℓ_1 and ℓ_2 norms. Finally, $\mathcal{L}_r(\cdot,\cdot)$ is the ℓ_2-loss between the sum of predicted outputs and the given input:

$$\mathcal{L}_r\left(\hat{\boldsymbol{\mu_d}}, \boldsymbol{\mu_T}\right) = \|\sum_d \hat{\boldsymbol{\mu_d}} - \boldsymbol{\mu_T}\|_2 \tag{13}$$

The training set is generated by first cropping CT volumes into desired sub-volumes. DRRs are generated from the original CT volume (stored in $\boldsymbol{\mu_T}$) as well as the sub-volumes (stored in $\boldsymbol{\mu_d}$) with fixed intrinsics parameters (estimated from the C-arm calibration), and different extrinsics parameters which simulate realistic trajectories of a C-arm.

Network Architecture: Our network architecture (c.f. Fig. 3) is similar to U-Net [10] with slight modifications, *i.e.* the last convolutional layers were dropped, and a dropout is added immediately after the last convolutional layer in the encoder to avoid overfitting and speed up the training process. In addition, we incorporate decomposition loss $\mathcal{L}_d$, which employs elastic-net loss, right after the last convolutional layer in the decoder part. To fulfill the reconstruction constraint (Beer-Lambert law), a 1×1 convolutional layer is added right before the reconstruction loss $\mathcal{L}_r$. To compensate for potential (rounding) errors due to normalization, interpolation, or the limited dynamic range we use a relaxed version for attenuation accumulation, where the reconstruced image is represented by a weighted sum of decomposed layers.

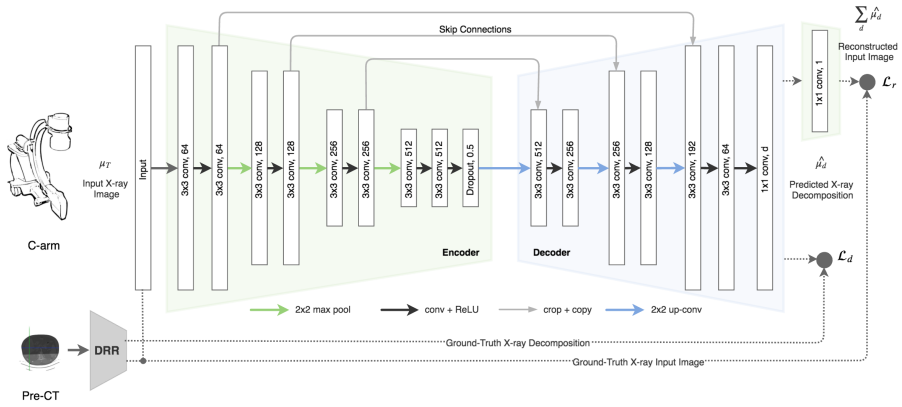

Fig. 3. X-ray decomposition framework: In training phase (dotted lines), DRRs are generated from preoperative CT scans and their corresponding clipped subvolumes, while in the testing phase, an X-ray image is acquired using C-arm and fed to the network to output the corresponding X-ray decompositions.

3 Experiments and Results

We have designed two experimintal setups for validating our proposed methodology. While the first experiment only targets patient-specific models, *i.e.* training and testing on the same patient, the second experiment validates the application-specific models, for *i)* intra-operative purposes, and *ii)* for diagnostic purposes, where Leave One Patient Out Cross Validation (LOPOCV) is performed for the latter. Models that are trained without the reconstruction loss ($\lambda_r = 0$) are considered as baseline models.

Dataset: We validated our proposed method on 6 clinical datasets (c.f. Table 1) acquired with different protocols for thorax and abdomen. Clipping planes are carefully selected to visualize anatomical regions of clinical interest such as vascualture tree, bifurcations, and specific anomalies (*e.g.* AAA shown in Fig. 1). For each clinicial dataset, approximately 1200 ground-truth DRRs are rendered for the whole volume as well as the corresponding sub-volumes, each of size 256×256 pixels (as described in Sec. 3). The generated DRRs are divided into non-overlapping sets of training (60%), validation (20%), and testing (20%). Patient-wise splitting is considered in some experimental setups, *i.e.* training (67%) and testing(33%).

Simulated X-ray of Clipped CT Volume: As depicted in Fig. 2, the pre-operative CT is clipped at two coronal planes of thorax (one coronal plane for abdomen) resulting in three sub-volumes; ribcage, vasculature and spine (two sub-volumes; AAA and spine for abdomen). These sub-volumes are padded with background values to retain the same dimension of the original CT. A C-arm simulator

Patient №	Volume (voxels)	Resolution (mm)	Protocol	Clipping plane @	source
1	$512 \times 482 \times 339$	$1 \times 1 \times 1$	Thorax	$175, 310$	TCIA
2	$350 \times 283 \times 367$	$1 \times 1 \times 1$	Thorax	$108, 175$	TCIA
4	$512 \times 386 \times 407$	$1 \times 1 \times 1$	Thorax	$164, 258$	TCIA
25	$512 \times 512 \times 175$	$0.92 \times 0.92 \times 3$	Abdomen	322	in-house
36	$512 \times 512 \times 517$	$0.92 \times 0.74 \times 1$	Abdomen	265	in-house
37	$512 \times 512 \times 144$	$0.85 \times 0.85 \times 3$	Abdomen	298	in-house

Table 1. Clinical datasets (TCIA courtesy of `www.cancerimagingarchive.net`)

software is used to generate DRRs along pre-defined trajectories, *i.e.* Cranial $(0 - 20°)$, Left Anterior Oblique (LAO) and Right Anterior Oblique (RAO) $(0 - 40°)$, for the original volume and different sub-volumes as well.

Implementation: The network archiecture together with the decomposition (elastic-net) and reconstruction losses are implemented in MATLAB and MatConvNet [11]. Network parameters, *i.e.* learning rate was set at 10^{-6}, the momentum at 0.9 and the batchsize at 16. The hyper-parameters λ_r was set at 0.5, where λ_d was altered between 0.1 and 0.9 to observe the influence of each term.

Evaluation metrics: We compute two different metrics to evalauate our results compared to the available ground truth; PSNR to evaluete the quality of reconstruction, and SSIM to measure the quality perception.

3.1 Patient-Specific Model

In this experimental setup, we train a baseline model on one patient and perform testing on the same patient for entirely different trajectories with additional translations. Quantitative results (c.f. Table 2) in terms of PSNR, confirm the visual results obtained in Fig. 4. It should be noted that this is the only model trained on pre-processed data, *i.e.* X-ray windowing. In the following section, we compare our proposed model with the baseline ($\lambda_r = 0$), then we take a step further and build more realistic scenarios to evaluate our proposed method.

Patient №	PSNR (dB)				SSIM (100%)			
	RibCage	Vascualture	Spine	Thorax	RibCage	Vascualture	Spine	Thorax
Overall	57.32	82.56	83.07	29.60	93.60	88.67	67.48	70.93

Table 2. Patient-Specific Model

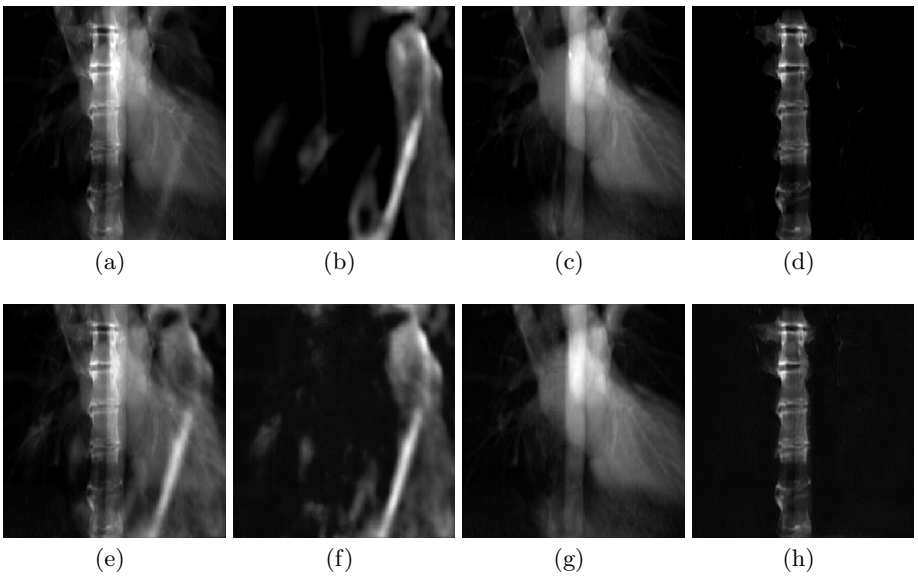

Fig. 4. Prediction of decomposed X-ray layers: (a) thorax X-ray, (b-d) ground-truth X-ray decompositions to ribcage, vasculature tree, and spine respectively, where regions of clinical interests are relatively obvious. (f-h) Predicted X-ray decomposition for a given input (a), where (e) the sum of (f-h) produces high reconstruction error.

3.2 Application-Specific Model

To address more realistic scenarios, we designed two experiments for: *i)* intra-operative scenarios, where the model is trained on several pre-operative patient CT data, and tested on the same patient data during the intervention, and *ii)* diagnostic applications which relies solely on X-ray image (patient CT not needed).

Intra-operative Purpose: In this setup, a pre-operative CT scan is required for a new patient, and is used to further fine-tune the pre-trained model. A model is trained on three patients (2160 images) for 200 Epochs and tested on the same patients for entirely different trajectories (720 images). Results are reported for both Abdomen and Thorax datasets in Table 3 and 4 respectively. A comparsion with baseline models ($\lambda_r = 0$) showing a positive infleunce of the reconstruction loss on the desired decompositions in terms of better PSNR and SSIM as well.

Diagnostic Purpose: In this setup, a preoperative CT scan is not required for the testing patient anymore. For this purpose, a model is trained on two patients (2400 images) for 60 Epochs ($\lambda_r = 0.5$) and tested on the third patient (1200 images). Leave One Patient Out Cross Validation (LOPOCV) on both Abdomen and Thorax datasets (c.f. Table 5 and 6) was performed to investigate

	Patient №	PSNR (dB)			SSIM (100%)		
		AAA	Spine	Abdomen	AAA	Spine	Abdomen
$\lambda_r = 0.5$	25	53.98	51.31	23.71	83.69	83.46	82.98
	36	56.08	55.31	27.08	87.22	86.17	81.96
	37	53.98	49.43	27.70	87.52	82.08	81.00
	Overall	**53.02**	**50.62**	26.00	**86.14**	**83.90**	**81.98**
$\lambda_r = 0$	25	51.56	50.89	23.58	81.05	83.02	82.16
	36	55.59	52.81	27.21	86.08	85.95	80.81
	37	52.87	48.79	27.96	86.99	82.04	80.10
	Overall	52.06	49.77	**26.05**	84.70	83.67	81.02

Table 3. Intra-operative purpose (Abdomen)

the robustness and validate the influence of hyper-parameter λ_d. For instance, high contribution of sparsity term ℓ_1 gives better results on thorax-spine image (SSIM=93.57) than the same contribution of smoothness term ℓ_2 (SSIM=91.90).

4 Discussion

This work presents a proof-of-concept of using machine learning for in-depth decomposition of single-view X-ray images. The methodology is evaluated on 6 medical datasets from chest and abdomen, and the results show that deep neural networks offer a promising outcome for this purpose. In the following, we discuss the statistics and our observations regrading the results presented in Sec. 3.

Ranges and X-ray Windowing: It came to our attention that X-ray image windowing by means of a pre-processing step has a positive influence on the training performance and results in faster convergance as well as a lower recontrucction error. This is done by suppressing the background intensities to zero, hence it

	Patient №	PSNR (dB)				SSIM (100%)			
		RibCage	Vascualture	Spine	Thorax	RibCage	Vascualture	Spine	Thorax
$\lambda_r = 0.5$	1	83.97	88.72	73.21	54.60	90.96	96.67	96.14	92.60
	2	74.74	84.54	78.85	57.21	89.54	93.22	95.42	92.88
	4	83.45	91.39	84.65	55.70	92.71	97.20	97.73	93.32
	Overall	**78.97**	87.36	**77.48**	**55.77**	**91.07**	**95.70**	**96.43**	**92.93**
$\lambda_r = 0$	1	70.41	80.30	65.56	54.45	80.55	94.23	93.34	91.63
	2	68.84	81.22	73.29	55.06	81.47	88.68	93.45	90.51
	4	72.77	87.18	76.23	54.03	85.98	94.43	95.78	91.59
	Overall	75.16	**87.44**	74.28	54.83	87.73	95.12	95.49	91.49

Table 4. Intra-operative purpose (Thorax)

	Patient №	PSNR (dB)			SSIM (100%)		
	train (test)	AAA	Spine	Abdomen	AAA	Spine	Abdomen
$\lambda_d = 0.1$	36,37 (25)	45.78	46.55	24.81	82.24	83.75	83.58
	25,37 (36)	50.78	50.40	25.08	83.76	83.21	79.54
	25,36 (37)	49.19	41.59	29.98	85.65	80.04	82.72
	Overall	48.36	**45.52**	26.36	83.89	**82.33**	81.95
$\lambda_d = 0.9$	36,37 (25)	45.84	46.56	24.78	82.26	83.81	83.73
	25,37 (36)	52.37	47.75	26.36	84.65	84.33	80.66
	25,36 (37)	48.15	40.21	30.85	85.60	77.80	83.32
	Overall	**48.44**	44.26	**27.02**	**84.17**	81.98	**82.57**

Table 5. LOPOCV for diagnostic purpose (Abdomen)

allows the network to better model the structures of interest such as rib cage, vasculature, and spine (c.f. Fig 4). In the contrary, using raw images results in learning additional background structures that demand a high model complexity. Lastly, we did not account for domain shift in the X-ray images used in our evaluation. Adoption of this method to real scenario X-ray imaging requires further investigation.

Independence Assumption: Ideally, the weights associated with last convolutional layer, immediately before the decomposition loss, should be orthogonal to each other. To investigate this, we run the orthogonality test, *i.e.* $\frac{1}{d}\sum_{d,i\neq j} 1 - (w_i^T \cdot w_j)$, and the average values are reported as approximately 0.51 and 0.61 for abdomen and thorax datasets, respectively. Imposing this constraint explicitly in the objective function is a potential future work.

Prior Knowledge: With the recent advances in X-ray imaging technology, phase contrast X-ray imaging has become an important topic of research [4], where

	Patient №	PSNR (dB)				SSIM (100%)			
	train (test)	RibCage	Vascualture	Spine	Thorax	RibCage	Vascualture	Spine	Thorax
$\lambda_d = 0.1$	2,4 (1)	73.23	80.28	61.70	60.62	83.19	95.53	90.83	95.16
	1,4 (2)	71.15	85.15	72.36	54.65	84.06	92.16	92.96	90.10
	1,2 (4)	78.19	92.42	78.21	53.79	89.13	96.61	96.93	92.09
	Overall	**73.78**	**84.78**	**68.40**	**55.93**	85.46	**94.10**	**93.57**	**92.45**
$\lambda_d = 0.9$	2,4 (1)	73.80	82.29	61.72	60.67	86.16	93.94	90.46	95.53
	1,4 (2)	67.48	84.48	66.01	54.30	84.38	92.44	90.88	90.49
	1,2 (4)	75.65	87.22	70.52	52.27	86.13	94.79	93.46	90.85
	Overall	71.67	84.46	65.45	55.16	**85.56**	93.72	91.90	92.29

Table 6. LOPOCV for diagnostic purpose (Thorax)

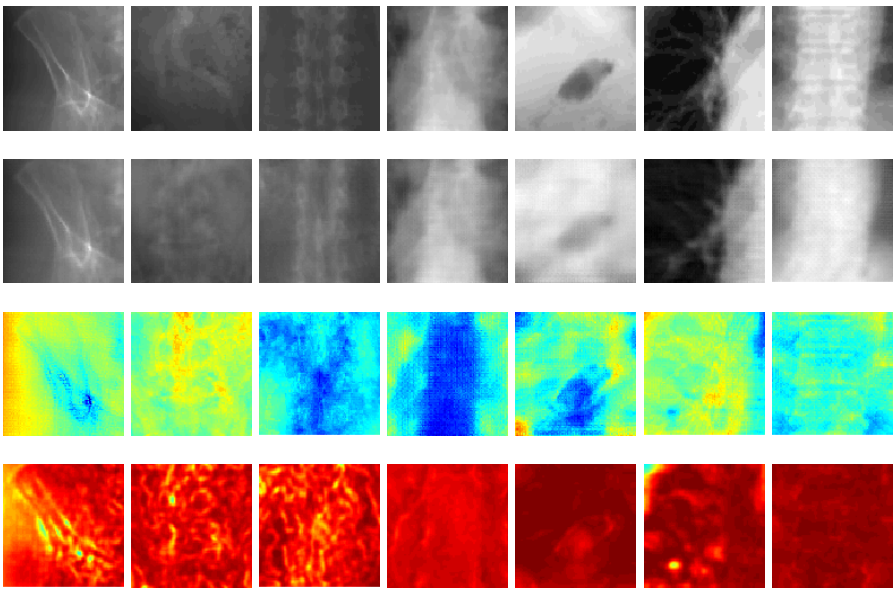

Fig. 5. Cropped ground-truth regions (row 1) of abdomen, AAA, and spine (column 1-3), thorax, ribcage, vasculature and spine (column 4-7). Predicted decompositions (row 2), residual and SSIM maps are visualized in jet colormap (row 3-4).

the X-ray phase, in addition to the X-ray attenuation, are obtained. We believe such information are highly valuable for our proposed method, and can serve as important prior knowledge.

Generalizability: Our results demonstrate promising results towards generalizability, in particular for diagnostic purposes. For instance, an implanted metal plate was present in patient 36 and recovered using a model trained on normal patients (c.f. Fig 6).

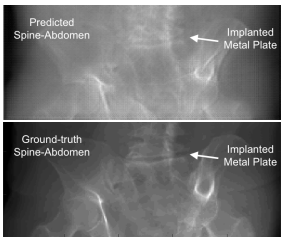

Fig. 6. X-ray pathology

Objective Function: In our methodology, we employed elastic-net, *i.e.* weighted sum of ℓ_1 and ℓ_2 norms, and investigated their contributions by varying the hyper-parameter λ_d to the lower and upper bounds. Results show a high correlation between this parameter and the desired channel, therefore we beleive this parameter should be chosen channel-wise.

We believe the in-depth X-ray decomposition presented in this work can improve the perceptual problems in standard X-ray radiographs, as well as in C-arm X-ray images for fluoroscopic and orthopedic interventions. This work reports very promising results, particularly for the visualization of obscured anatomy. Yet, additional clinical validation is required for diagnostic and interventional setups.

References

1. Albarqouni, S., Konrad, U., Wang, L., Navab, N., Demirci, S.: Single-view x-ray depth recovery: toward a novel concept for image-guided interventions. International journal of computer assisted radiology and surgery pp. 1–8 (2016)
2. Bravin, A., Coan, P., Suortti, P.: X-ray phase-contrast imaging: from pre-clinical applications towards clinics. Physics in medicine and biology 58(1), R1 (2012)
3. Krupinski, E.A.: The importance of perception research in medical imaging. Radiation medicine 18(6), 329–334 (2000)
4. Lewis, R.: Medical phase contrast x-ray imaging: current status and future prospects. Physics in medicine and biology 49(16), 3573 (2004)
5. Mazess, R.B., Barden, H.S., Bisek, J.P., Hanson, J.: Dual-energy x-ray absorptiometry for total-body and regional bone-mineral and soft-tissue composition. The American journal of clinical nutrition 51(6), 1106–1112 (1990)
6. Miao, S., Wang, Z.J., Liao, R.: A cnn regression approach for real-time 2d/3d registration. IEEE transactions on medical imaging 35(5), 1352–1363 (2016)
7. Mistretta, C.A., Crummy, A.B., Strother, C.M.: Digital angiography: a perspective. Radiology 139(2), 273–276 (1981)
8. Robinson, P.: Radiology's achilles' heel: error and variation in the interpretation of the röntgen image. The British Journal of Radiology 70(839), 1085–1098 (1997)
9. Robinson, P., Wilson, D., Coral, A., Murphy, A., Verow, P.: Variation between experienced observers in the interpretation of accident and emergency radiographs. The British journal of radiology 72(856), 323–330 (1999)
10. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 234–241. Springer (2015)
11. Vedaldi, A., Lenc, K.: Matconvnet: Convolutional neural networks for matlab. In: Proceedings of the 23rd ACM international conference on Multimedia. pp. 689–692. ACM (2015)
12. Wang, J., Fallavollita, P., Wang, L., Kreiser, M., Navab, N.: Augmented reality during angiography: integration of a virtual mirror for improved 2d/3d visualization. In: Mixed and Augmented Reality (ISMAR), 2012 IEEE International Symposium on. pp. 257–264. IEEE (2012)
13. Wang, J., Kreiser, M., Wang, L., Navab, N., Fallavollita, P.: Augmented depth perception visualization in 2d/3d image fusion. Computerized Medical Imaging and Graphics 38(8), 744–752 (2014)
14. Wang, X., zu Berge, C.S., Demirci, S., Fallavollita, P., Navab, N.: Improved interventional x-ray appearance. In: Mixed and Augmented Reality (ISMAR), 2014 IEEE International Symposium on. pp. 237–242. IEEE (2014)
15. Wieczorek, M., Aichert, A., Fallavollita, P., Kutter, O., Ahmadi, A., Wang, L., Navab, N.: Interactive 3d visualization of a single-view x-ray image. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 73–80. Springer (2011)
16. Zou, H., Hastie, T.: Regularization and variable selection via the elastic net. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 67(2), 301–320 (2005)